

10/04/2004

L Number	Hits	Search Text	DB	Time stamp
1	48843	router	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 09:42
2	3308	internal\$6 with (backup or back-up or (back adj1 up))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:10
3	16	router with (internal\$6 with (backup or back-up or (back adj1 up)))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:05
4	5768	switch\$6 adj3 fabric	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:10
5	84	(backup or back-up or (back adj1 up)) adj1 router	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:13
6	4	(switch\$6 adj3 fabric) same ((backup or back-up or (back adj1 up)) adj1 router)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:31
7	2	("20020167952").PN.	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:32
8	0	("20020167952").PN.) and (switch\$6 adj3 fabric)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:32
9	32058	switchover or switch-over or (switch adj1 over)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:33
10	2820	failover or fail-over or (fail adj1 over)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:36

11	34624	((switchover or switch-over or (switch adj1 over)) or (failover or fail-over or (fail adj1 over))	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:36
12	8	((switchover or switch-over or (switch adj1 over)) or (failover or fail-over or (fail adj1 over))) adj1 router	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:45
13	0	((switchover or switch-over or (switch adj1 over)) or (failover or fail-over or (fail adj1 over))) adj1 router) and (switch\$6 adj3 fabric)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:37
14	38	((switchover or switch-over or (switch adj1 over)) or (failover or fail-over or (fail adj1 over))) adj3 router	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:45
15	110	((backup or back-up or (back adj1 up)) adj1 router) or ((switchover or switch-over or (switch adj1 over)) or (failover or fail-over or (fail adj1 over))) adj3 router)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 10:46
16	14	(((backup or back-up or (back adj1 up)) adj1 router) or ((switchover or switch-over or (switch adj1 over)) or (failover or fail-over or (fail adj1 over))) adj3 router)) and (switch\$6 adj3 fabric)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 11:01
17	25486	(primary or second\$4 or redundan\$4 or two or backup or back-up) adj2 processor\$	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 11:05
18	7	((primary or second\$4 or redundan\$4 or two or backup or back-up) adj2 processor\$) and (((backup or back-up or (back adj1 up)) adj1 router) or ((switchover or switch-over or (switch adj1 over)) or (failover or fail-over or (fail adj1 over))) adj3 router)) and (switch\$6 adj3 fabric)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 11:07
19	954	primary adj1 port	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 11:12
20	2	((primary or second\$4 or redundan\$4 or two or backup or back-up) adj2 processor\$) and (((backup or back-up or (back adj1 up)) adj1 router) or ((switchover or switch-over or (switch adj1 over)) or (failover or fail-over or (fail adj1 over))) adj3 router)) and (switch\$6 adj3 fabric)) and (primary adj1 port)	USPAT; US-PGPUB; EPO; JPO; DERWENT; IBM_TDB	2004/05/24 11:12

DOCUMENT-IDENTIFIER: US 20020060986 A1

TITLE: ROUTER DEVICE HAVING A REDUNDANT CONFIGURATION

----- KWIC -----

Summary of Invention Paragraph - BSTX (21):

[0020] To prevent the other routers from being affected by this system switchover of the route calculation units or to facilitate the system switchover, a thinkable method includes the steps of sending all items of information (network link-state information, states of routers, and states of interfaces), which the route calculation unit in the active mode obtained from the routing protocol process, from the route calculation unit in the active mode to the route calculation units in the standby mode, and storing them in the route calculation units in the standby mode. With the above arrangement, the same states in the route calculation unit, which was previously in the active mode, can be reproduced in a route calculation unit which is subsequently brought into the active mode, and the route calculation unit newly brought into the active mode can promptly become capable of executing the same process as did the previously operating route calculation unit. Consequently, the other routers are protected from affects of the system switchover of the routers.

US-PAT-NO: 6618389

DOCUMENT-IDENTIFIER: US 6618389 B2

TITLE: Validation of call processing network
performance

----- KWIC -----

Detailed Description Text - DETX (82):

In an example failover test run, for example, a UDP packet stream is sent from the bliss01 server round trip to a distant system via the Cisco 7513 routers (FIG. 3). The routers were set up using HSRP as the failover mechanism. The primary router was then powered off. The characteristics of the failovers on the routers due to power loss include: calculation of event duration from the start of the buffer loss, e.g., at one millisecond per buffer, to the end of the buffer delay period, where the delay returns to the nominal delay value.

US-PAT-NO: 6490246

DOCUMENT-IDENTIFIER: US 6490246 B2

TITLE: System and method for using active and standby
routers
wherein both routers have the same ID even before
a
failure occurs

----- KWIC -----

Brief Summary Text - BSTX (21):

To prevent the other routers from being affected by this system switchover
of the route calculation units or to facilitate the system switchover,
a
thinkable method includes the steps of sending all items of information
(network link-state information, states of routers, and states of
interfaces),
which the route calculation unit in the active mode obtained from the
routing
protocol process, from the route calculation unit in the active mode to
the
route calculation units in the standby mode, and storing them in the
route
calculation units in the standby mode. With the above arrangement, the
same
states in the route calculation unit, which was previously in the
active mode,
can be reproduced in a route calculation unit which is subsequently
brought
into the active mode, and the route calculation unit newly brought into
the
active mode can promptly become capable of executing the same process
as did
the previously operating route calculation unit. Consequently, the
other
routers are protected from the affects of the system switchover of the
routers.

DOCUMENT-IDENTIFIER: US 20030237016 A1

TITLE: System and apparatus for accelerating content delivery throughout networks

----- KWIC -----

Detail Description Paragraph - DETX (25):

[0056] XP 548 is generally responsible for managing NPs 518a-518d as well as coordinating the functions of NPs 518a-518d with any external processors such as CPU 503. FP 551 is generally responsible for scaling network processors 518a-518d with switching fabrics such as packet routing switch IC 533. TLU 554 is generally responsible for implementing table searches and updates and QMU 557 is generally responsible for integrating queue control and management. BMU 560 is generally responsible for providing fast, flexible memory management.

Detail Description Paragraph - DETX (46):

[0077] Host processor subsystem 706 preferably includes the following functionality: initiates system boot and initialization sequence for content router 200; provides an external interface to serial port 318 for local and remote systems management; implements systems management functionality, including SNMP, Web based management interfaces, and serial port interface; handshakes with rapidstack subsystem 703 to collect data to support systems management; provides systems management platform APIs that can be exposed to implement a custom systems management presentation layer; accepts redirected file writes from rapidstack subsystem 703, implemented as a network attached storage (NAS) device supporting common internet file system (CIFS) and network file system (NFS) protocols; handshakes with rapidstack subsystem 703 to synchronize read/write access to local storage that flows through host processor 706; handshakes with rapidstack subsystem 703 to maintain a current file system directory; handshakes with rapidstack subsystem 703 to maintain cache coherency and implements clustering features such as failover from one

content router 200 to another content router 200 in a cluster and auto-discovery and storage replication to each content router 200 added to such a cluster.

Detail Description Paragraph - DETX (93):

[0124] With respect to content router 200 to content router 200 connectivity in particular, control functions and systems management data preferably flow between the content routers 200 on cluster LAN 1005 in order for the content routers 200 to operate as a cluster. Examples of such data may include heart beat packets facilitating failover from one content router 200 to another, the gathering of performance data by a primary content router 200 from other content routers 200 in the cluster as well as other data.

Detail Description Paragraph - DETX (110):

[0141] To provide reliable service, content router 200 is preferably operable to provide high levels of fault tolerance. Content routers 200 existing within a cluster preferably include the ability to provide fault tolerance through failover from a primary content router 200 to at least a secondary content router 200. As such, the minimum fault tolerance functionality preferred is the ability to failover from a primary content router to a redundant "hot spare" content router 200. Client systems coupled to a primary content router 200 that fails may experience loss of connection. However, the clients are generally capable of re-coupling quickly.

DOCUMENT-IDENTIFIER: US 20030097428 A1

TITLE: Internet server appliance platform with
flexible integrated suite of server resources and content
flow delivery capabilities supporting continuous data
demands and bursty demands

----- KWIC -----

Summary of Invention Paragraph - BSTX (13):

[0007] The physical environment in which networked computing equipment is housed is usually very "computer" oriented, wherein the systems employed must be completely shut down for maintenance or service such as replacing a circuit card or performing loading of software. Most of these systems provide no redundancy within the computers themselves, so redundancy must be achieved by having duplicate computing units interconnected to the local data networks. To switch over from use of one computer to another, a router change is made such that all new "sessions" of applications are directed towards the back up computer, which eventually frees up the primary computer so that it can be shut down and serviced.

Detail Description Paragraph - DETX (8):

[0039] In this inventive arrangement, processing resources and storage resources are decoupled from each other, both physically and logically.

Additionally, a switching fabric is built into the architecture of the system. In this modular organization, a very wide variety of configurations of the Internet Server Appliance Platform (ISAP) system may be realized by installing more or less of each resource type, and soft reconfiguring the system to utilize those resources per the requirements of specific application programs.

Detail Description Paragraph - DETX (17):

[0048] Due to the system's built-in switching fabric with integrated and scalable processing and storage facilities, considerable cable bulk is

eliminated which is normally present in systems comprised of multiple racks and individual computing units. This reduces a typical cable harness diameter of 5.6 inches to approximately 1.5 inches for 200 server units. This reduces the cost of ownership of the system by increasing reliability, improving maintainability, and reducing sheer bulk and space requirements. Top Level Architecture

DOCUMENT-IDENTIFIER: US 20020099972 A1

TITLE: Method and system for selecting a master
controller in
a redundant control plane having plural
controllers

----- KWIC -----

Summary of Invention Paragraph - BSTX (6):

[0005] Typically, network devices include line cards that have input and output ports coupled to communication links. The line cards are also typically coupled to a switch fabric. Data units received at an input port are forwarded to the switch fabric and to at least one output port of at least one line card for forwarding over the corresponding communication link.

Detail Description Paragraph - DETX (3):

[0026] Consistent with the present invention, a redundant control plane architecture for use in a network device such as a network switch or network router is disclosed. The redundant control plane permits failover or a controlled switchover from a Master or active Router Control Processor (RCP) to a standby RCP in the event of a failure of any single component or bus within the redundant control plane so that the control plane remains functional. The control paths in the presently disclosed redundant control plane are isolated from the data plane of the network device so that control traffic does not utilize data plane bandwidth. The control plane in the presently disclosed system is used as the communication path for downloading forwarding tables to the line cards, for transmittal of control and configuration information to be stored in the line cards and for transmittal of slow path traffic from line cards to a central processing function for handling. Such slow path traffic includes, for example, Internet Protocol (IP) packets sourced by or addressed to a virtual router, IP packets requiring IP Option processing, IP packets requiring IP fragmentation and Layer 2 control protocol communications.

Detail Description Paragraph - DETX (4):

[0027] Components included within the presently disclosed network device 100 are depicted in FIG. 1. The network device 100 includes a midplane 101 that selectively interconnects a plurality of printed circuit cards within the network device. More specifically, the printed circuit cards illustrated in FIG. 1 include first and second Router Control Processors (RCPs) 102a, 102b, first and second Bridge Hotswap Cards (BHCs) 104a, 104b and first and second pluralities of line cards 108a and 108b. The line cards each typically include one or more input ports 113 for receiving data over input communication links 115 and one or more output ports 117 for forwarding data from the network device 100 over associated output communication links 119 (one each shown; see FIG. 2). The network device 100 further includes first and second Switch Cards 106a, 106b, first and second switch fabrics 111a and 111b (see FIG. 2) within the first and second Switch Cards 106a and 106b, respectively, physical device interfaces (PHYS) as known in the art (not shown) for connection of the Line Cards 108a, 108b to the applicable communication media, and a common I/O card 103 (see FIG. 2) that includes logic used in the selection of a Master RCP/BHC pair. The Line Cards 108a, 108b, the Bridge Hotswap cards 104a, 104b and the Switch Cards 106a, 106b are electrically interconnected as discussed below via conductive traces on the midplane 101. The RCP cards 102a, 102b are communicably coupled to the Bridge Hotswap cards 104a and 104b respectively, and reside in the rear portion of the midplane slots occupied by the Bridge Hotswap cards 104a and 104b respectively.

Detail Description Paragraph - DETX (10):

[0033] As indicated above, the Switch Card A 106a and Switch Card B 106b are the primary conduits for interRCP communications and additionally include switch fabrics 111a and 111b respectively for performing data forwarding within the data plane. Either Switch Card A 106a or Switch Card B 106b may be the active Switch Card for inter-RCP communication at any given time and the other Switch Card is the standby Switch Card for interRCP communication during normal

operation.

Detail Description Paragraph - DETX (22):

[0045] FIG. 4 depicts an illustrative block diagram of the Switch Card 106 that corresponds to the Switch Card A 106a and Switch Card B 106b depicted in FIGS. 1 and 2. The block diagram only depicts the portion of the Switch Card logic that pertains to the redundant control plane and does not illustrate the switch fabric that is employed within the data plane for traffic forwarding. The Switch Card A 106a and the Switch Card B 106b comprise the primary pathway for inter-RCP communications. The RCP-to-RCP communications are used primarily to maintain consistency between the RCP master and RCP standby device. RCP state information is passed between RCPs using one of the Switch Cards 106a, 106b as a conduit. The Switch Card that serves as the conduit for transfer of information may be either the Switch Card on which the active switch fabric resides or the other Switch Card.

Detail Description Paragraph - DETX (24):

[0047] Each Switch Card 106 also includes control status registers (CSRs) 306 that hold control information for the respective Switch Card. For example, the CSRs contain a switch fabric reset bit, an HDLC controller enable bits and interrupt mask registers. The CSR registers also contain status information for the respective Switch Card 106 that identifies whether the Switch Card is the active Switch Card, Switch identifier, a revision identifier, interrupt and error bits.

Detail Description Paragraph - DETX (25):

[0048] The HDLC Controller 304 is used to communicate with a processor within the Switch Fabric Subsystem 308 on the respective Switch Card 106. All communication with the processor and the RCP is via HDLC commands.

Claims Text - CLTX (2):

1. Apparatus for configuring a control plane in a network device having at least one line card for receiving and transmitting data, said at least one line

card being communicably coupled to at least one switch fabric via at least one data path, said apparatus comprising: first and second control plane processor elements, each operative to generate at least one status signal indicative of the operational status of the respective processor element; first and second control paths associated with said first and second control plane processor elements respectively and communicably coupling said first and second control plane processor elements respectively to said at least one line card, said first and second control paths being isolated from said at least one data path; selection logic operative in response to said at least one status signal from said first and second control plane processor elements for generating at least one identification signal for identifying one of said first and second control plane processor elements as an Master control plane processor and for communicating said at least one identification signal to said first and second control plane processor elements; one of said first and second control plane processor elements being operative in response to receipt of said at least one identification signal to configure itself as said active control plane processor element; and said active control plane processor element being operative to transmit configuration information over the associated control path to said at least one line card.

Claims Text - CLTX (11):

10. A method for configuring a control plane in a network device, wherein said network device includes at least one line card for receiving and transmitting data over corresponding communication links, said at least one line card being communicably coupled to at least one switch fabric via at least one data path, wherein said at least one line card, said at least one data path and said at least one switch fabric comprise a data plane, said method comprising the steps of: generating at least one status signal at first and second control plane processor elements, each of said at least one signal being indicative of the ability of the respective control plane processor element to function as an active control plane processor element for said control plane; communicating said at least one status signal from said first and second

control plane processor elements to selection logic; generating within said selection logic at least one identification signal responsive to at least one status signal from said first and second control plane processor elements, said identification signal for identifying one of said first and second control plane processor element as said active control plane processor element; communicating said at least one identification signal from said selection logic to said first and second control plane processor elements; responsive to receipt of said at least one identification signal at said first and second control plane processor elements, configuring one of said first and second control plane processor elements as an active control plane processor element and the other one of said processor elements as a standby control plane processor element; and transmitting first configuration information from said active control plane processor element to said at least one line card over a first control path isolated from said at least one data path.

DOCUMENT-IDENTIFIER: US 20010048661 A1

TITLE: Method and apparatus for multi-protocol
redundant
router protocol support

----- KWIC -----

Summary of Invention Paragraph - BSTX (4):

[0003] Redundant routing protocols have been developed to provide hosts configured with static routes a measure of protection against router failure. In redundant routing, a host is configured to send to a virtual router MAC address that is supported by two or more physical routers sharing a LAN with the host. Particularly, at any given time in an operational cycle, one of the physical routers, a virtual master, is responsible for forwarding packets received from the host and having the virtual router MAC address, and the other backup routers standby to assume forwarding responsibilities in the event the virtual master fails. The transition by which respective ones of the backup routers become the virtual master is transparent to the host.

Detail Description Paragraph - DETX (8):

[0023] In FIG. 1, the routers 110 and 116 are illustrated to be supporting four groups of virtual routers (i.e., HSRP Group 1 virtual router, VRRP Group 2 virtual router, HSRP Group 2 virtual router, and VRRP Group 1 virtual router). Therefore, for example, when the router 110 operates as the HSRP Group 1 virtual master and the VRRP Group 2 virtual master, the router 116 may operate as an HSRP Group 1 standby router and a VRRP Group 2 standby router. For another example, when the router 116 operates as the HSRP Group 2 virtual master and the VRRP Group 1 virtual master, the router 110 may operate as an HSRP Group 2 standby router and a VRRP Group 1 standby router. Standby routers may also be referred to as backup routers. In other embodiments, each physical router may be mapped to one HSRP group of redundant routers and one VRRP group of redundant routers.

Detail Description Paragraph - DETX (18):

[0033] The packet buffer 202 may also include an edit module for editing the packets prior to forwarding them out of the switching controller as outbound packets 218. The edit module may include an edit program construction engine for creating edit programs real-time and/or an edit engine for modifying the packets. The application engine 206 preferably provides application data 216, which may include a disposition decision for the packet, to the packet buffer 202, and the edit program construction engine preferably uses the application data to create the edit programs. The outbound packets 218 may be transmitted over a switching fabric interface to communication networks, such as, for example, the Ethernet.

DOCUMENT-IDENTIFIER: US 20030218982 A1

TITLE: Highly-available OSPF routing protocol

----- KWIC -----

Summary of Invention Paragraph - BSTX (10):

[0008] The present invention is directed to a system and method of highly-available Open Shortest Path First (OSPF) routing in a network. The dynamic state of a backup OSPF instance in a router is synchronized with the dynamic state of an active OSPF instance using explicit message transmission from the active instance to the backup instance in the router control plane. After this the dynamic state synchronization of the backup OSPF instance is maintained using a combination of explicit message updates from the active OSPF instance together with a message flow-through mechanism. The active OSPF maintains forwarding tables in a shared central data plane that routes transit traffic through a shared central switch fabric. In the event of failure of the active OSPF instance, then the router recovers seamlessly without reconfiguring or interrupting traffic among peer routers in the network, by functionally substituting the synchronized backup OSPF instance for the active OSPF instance, such that the backup OSPF instance establishes itself as the new active OSPF instance. During this recovery process, the shared central switch fabric in the shared central data plane continues to forward transit traffic in accordance with route instructions implemented through forwarding tables created and maintained by the control plane.

Detail Description Paragraph - DETX (2):

[0015] This invention is directed to a novel architecture for high-availability (HA) OSPF dynamic routing protocol. The purpose of HA OSPF is to provide fail-over protection in the event of failure or shutdown of the hardware platform supporting the routing protocol control plane. Fail-over protection means that the OSPF function of the control plane continues to operate and maintains all dynamic state information. OSPF instances on other

routers in the network do not detect the fail-over condition on the local router. The HA OSPF architecture is based on commodity hardware components and is completely compatible with existing OSPF standards and available implementations. The architecture uses two separate hardware platforms that each execute the control plane software. The two systems are assigned roles of "active" and "backup." The active and backup protocol processors are connected to each other by a special purpose network.

Detail Description Paragraph - DETX (3):

[0016] Dynamic routing protocols are protocols that routers use to communicate with each other, to decide where the traffic goes on the Internet. In "Highly available (HA) routing protocols", routing fails over completely seamlessly. The outside world is unaware that there has been a fault from one router to another. The backup software and the backup router take over seamlessly, such that no one in the outside world knows that there has been a problem. During this recovery process, a central switch fabric in the central data plane of the router continues to forward transit traffic in accordance with routing instructions in forwarding tables created and maintained by the control plane.

Detail Description Paragraph - DETX (7):

[0020] For protection at the OSPF system level and the BGP level, the starting point after the communication mechanisms between routers are protected is to make sure that these can be recovered during switch-over from an active protocol processor to a backup protocol processor, to protect the actual software running on a high level, for example, OSPF or BGP, such that all the state information and detailed operations that are running BGP or OSPF are successfully and seamlessly transported to the backup system, thus allowing the backup system to take over. A number of algorithms and procedures are executed to accomplish that seamless fail-over.

Detail Description Paragraph - DETX (11):

[0024] FIG. 2A schematically illustrates router subsystems 20 of the HA OSPF architecture. Control plane 212 includes active protocol processor 21

running
active OSPF software instance 23 and backup protocol processor 22
running
backup OSPF software instance 24. Active and backup OSPF protocol
instances
are directly linked together through reliable TCP MNET 201. Routing
protocol
control traffic to and from peer OSPF routers 10-2, . . . , 10-N is
distributed within control plane 212 through routing network (RNET)
202. HA
OSPF features a flow-through architecture, such that all incoming
protocol
control traffic flows first via RNET 202 over link 204 through backup
instance
24 before flowing from backup instance 24 over link 205 to active
instance 22.
Similarly, all protocol control traffic originating at active instance
22 flows
first over link 206 through backup instance 24 before flowing from
backup
instance 24 over link 207 into RNET 202. From RNET 202 the protocol
control
traffic is distributed to peer OSPF routers 10-2, . . . , 10-N.

Detail Description Paragraph - DETX (12):

[0025] Active OSPF instance 22 alone executes the SPF algorithm
using input
protocol control information to initialize and update forwarding tables
25-1, . . . , 25-N in shared central data plane 211. Transit traffic to and
from peer
OSPF routers 10-2, . . . , 10-N travels through data links 203-1, . . . ,
203-N and is routed by shared central multipole optical switch fabric
27 within
shared central data plane 211 in accordance with forwarding tables
25-1, . . . , 25-N through respective associated packet forwarding engines (PFEs)
26-1, . . . , 26-N. Control and transit traffic to and from other peer OSPF
routers
(not shown) flows through router network interfaces 214.

Detail Description Paragraph - DETX (13):

[0026] In the event that active protocol processor 21 fails or must
be shut
down to perform maintenance, backup protocol processor 22 assumes the
functions
required for control plane 212. Data plane 211 continues to forward
transit
packets using forwarding tables 25-1, . . . , 25-N and is unaware of
the
processor switch-over in control plane 212. Thus, the function of
backup
protocol processor 22 is to maintain sufficient static and dynamic
state, so

that it can assume the role of active protocol processor 21 at any time.

Detecting the failure of the active protocol processor is performed by the operating system.

Detail Description Paragraph - DETX (14):

[0027] In the HA architecture, OSPF on backup protocol processor 22 maintains state information corresponding to the OSPF state of active protocol processor 21 with regard to neighbor relationships/adjacencies and link state database. As such, OSPF on the backup protocol processor is a passive consumer of information originated by active OSPF instance 23 and other OSPF routers in the network. Backup OSPF instance 24 does not send any protocol packets nor does it execute the SPF algorithm. It does not create any LSAs and does not execute any timer-driven functions. In the event of a fail-over, backup OSPF instance 24 executes recovery functions, such that it begins sending the same Hello packets as former active OSPF instance 23, assumes the LSA flooding functions of the former active instance, and executes the SPF algorithm to update forwarding tables 25-1, . . . , 25-N.

Detail Description Paragraph - DETX (34):

[0047] In the event of a fail-over at step 227, message flow-through ceases, such that the backup protocol processor communicates directly with other routers in the network. In HA OSPF architecture, the backup OSPF instance establishes itself as the active OSPF instance by enabling periodic timer processing to perform OSPF maintenance functions, for example, transmitting Hello packets and refreshing self-originated LSAs. During recovery, OSPF also executes the SPF algorithm on its existing link-state database. The forwarding table calculated from the SPF algorithm is then sent to shared central data plane 211 to be used in forwarding transit traffic. The assumption is that the backup OSPF instance has been able to maintain sufficient state synchronization with the active, so that when the fail-over occurs, the backup already has the data available in its memory to perform this processing.